

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Services and Communications Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>					
1. REPORT DATE (DD-MM-YYYY) 2 APR 07		2. REPORT TYPE FINAL REPORT		3. DATES COVERED (From - To) 15 MAY 01 - 30 APR 06	
4. TITLE AND SUBTITLE (CIP/SW-01) ARCHITECTURES FOR SECURE AND ROBUST DISTRIBUTED INFRASTRUCTURES				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER F49620-01-1-0365	
				5c. PROGRAM ELEMENT NUMBER 61103F	
6. AUTHOR(S) PROF SANJAY LALL				5d. PROJECT NUMBER 5094	
				5e. TASK NUMBER JS	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) STANFORD UNIVERSITY DEPT OF AERONAUTICS & ASTRONAUTICS 496 LOMITA MALL, DURAND BUILDING, STANFORD, CA 94305-4035				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NL 875 NORTH RANDOLPH STREET SUITE 325, ROOM 3112 ARLINGTON, VA 22203-1768				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVE FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>The major barrier constraining the successful management and design of large-scale distributed infrastructures is the conspicuous lack of knowledge about their dynamical features and behaviors. Up until very recently analysis of systems such as the Internet, or the national electricity distribution system, have primarily relied on the use of non-dynamical models, which neglect their complex, and frequently subtle, inherent dynamical properties. These traditional approaches have enjoyed considerable success while systems are run in predominantly cooperative environments, and provided that their performance boundaries are not approached. With the current proliferation of applications using and relying on such infrastructures, these infrastructures are becoming increasingly stressed, and as a result the incentives for malicious attacks are heightening. The stunning fact is that the fundamental assumptions under which are significant large-scale distributed infrastructures have been constructed and analyzed on longer hold; the invalidity of these non-dynamical assumptions is witnessed with the greater frequency of catastrophic failures in major infrastructures.</p>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)

Final Report: Architectures for Secure and Robust Distributed Infrastructures*

Sanjay Lall, Aeronautics and Astronautics, Stanford
Carolyn Beck, General Engineering, UIUC
Stephen Boyd, Electrical Engineering, Stanford
John Doyle, Control and Dynamical Systems, Caltech
Geir Dullerud, Mechanical Engineering, UIUC
Chris Hadjicostis, Electrical and Computer Engineering, UIUC
Muriel Medard, Electrical Engineering and Computer Science, MIT
Balaji Prabhakar, Electrical Engineering and Computer Science, Stanford
Rayadurgam Srikant, General Engineering UIUC
George Verghese, Electrical Engineering and Computer Science, MIT

Abstract

The major barrier constraining the successful management and design of large-scale distributed infrastructures is the conspicuous lack of knowledge about their *dynamical* features and behaviors. Up until very recently analysis of systems such as the Internet, or the national electricity distribution system, have primarily relied on the use of non-dynamical models, which neglect their complex, and frequently subtle, inherent dynamical properties. These traditional approaches have enjoyed considerable success while systems are run in predominantly cooperative environments, and provided that their performance boundaries are not approached. With the current proliferation of applications using and relying on such infrastructures, these infrastructures are becoming increasingly stressed, and as a result the incentives for malicious attacks are heightening. The stunning fact is that the fundamental assumptions under which all significant large-scale distributed infrastructures have been constructed and analyzed no longer hold; the invalidity of these non-dynamical assumptions is witnessed with the greater frequency of catastrophic failures in major infrastructures such as the Internet, the power grid, the air traffic system, and national-scale telecommunication systems.

This project is about network, reliability and robustness in large-scale systems. The major vision of this program is ubiquitous: we have distributed computing and information and would like to link these via secure communications to allow coordination of limited resources to achieve

*This research was supported by AFOSR DoD award number 49620-01-1-0365

global objectives that can be both predicted and guaranteed. The objective is to ensure that incorrect local decisions, due to dynamical effects, do not cause large-scale failures.

To address the challenges posed by dynamical behavior of large-scale network infrastructures, we bring to bear the tools and techniques of control theory together with those from communication networks and queuing theory. In particular, the algorithms and analytical approaches of control used for developing control strategies and logic are combined with protocol design methods to construct new, secure architectures for distributed networks. We focus on the dominant issues of complex dynamic behavior, local rather than global information and state, distributed rather than centralized decision making, secure, robust performance in an uncertain environment, and dynamic network connectivity.

Contents

1	Introduction	3
2	Congestion and Buffering in Wired networks	4
2.1	Sizing issues in buffer routers	5
2.2	Fluid model development and analysis of priority processing schemes in the Internet	6
2.3	Connection-level stability analysis in the Internet	7
3	Scheduling and Resource Allocation in Wireless Networks	7
3.1	A Large Deviations Analysis of Scheduling	7
3.2	Distributed Fair Resource Allocation in Cellular Networks in the Presence of Heterogeneous Delays	9
3.3	Simultaneous Routing and Resource Allocation	11
4	Fault Diagnosis over Packet Dropping Networks	11
4.1	Probabilistic Fault Detection and Identification	12
4.2	Distributed Symmetric Function Computation in Noisy Wireless Sensor Networks with Binary Data	13
5	Decentralized Control	15
5.1	Control over Networks	15
5.2	Decentralized Observation and Monitoring	16
5.3	Decentralized Control	17
5.4	Monitoring and Control of Power System Dynamics	18
6	Complexity and robustness in complex networks	18

7 Network Reliability	23
7.1 Graph Structure and Similarity	24
8 Impact of research and transitions	25

1 Introduction

The objective of this project is to advance the security and reliability of large-scale network infrastructures. The research is specifically targeted at the most critical medium- and long-term security and performance issues facing current and future military networks, as well as homeland installations. The program is focused both on fundamental scientific understanding of the mechanisms by which single attacks can lead to catastrophic cascading failures caused by dynamic effects of propagation, misprioritization and instability, as well as development of new protocols which eliminate such vulnerabilities. The research of this URI is already having significant impact on the commercial sector, affecting the router designs of Cisco, and the protocol designs of Microsoft and Nokia, as well as the open protocols behind TCP and AQM.

The program 1) analyzes the decisions made by routers and layer protocols to see how they lead to network-level consequences; 2) studies the propagation dynamics of the network to characterize instabilities; 3) has developed protocols that eliminate these instabilities; 4) studies the measurable indicators of high-throughput data streams that can be used to detect attempted attacks in real-time and monitor network performance; 5) formulates new methods, such as combining routing and coding over multiple paths, that make certain classes of attack much more difficult; 6) develops information-based randomized algorithms that prevent attacks which depend on distributed, multi-source simultaneous attack and response; 7) re-examines the basic protocols of the network to suggest modifications that can be incrementally deployed, without requiring all users to simultaneously change software systems; 8) understands the parallels and distinctions between data networks and transportation and energy networks, noting that each of these three infrastructures has witnessed large-scale catastrophic failures of similar nature in recent years; and 9) has worked on a fundamental rethinking of how such networks function, to guide designers of the next generation of systems.

Traditional approaches to increasing network security have primarily focused on protecting the integrity of nodes at the edge of the network rather than systematic and robust design of the network itself. In these traditional approaches the emphasis has been on improving encryption, key-exchange protocols and intelligent attack detection, so as to achieve nodes with extremely fortified defenses. This does not however protect against the catastrophic failures we have witnessed in recent years on the Internet and the National Power Grid, where nodes misinform each other, or under- or over-react to remote events, causing large-scale cascading failures. In current networks knowledge of which nodes to rely upon to make mission-critical decisions is being passed through and fed back

via increasingly long chains of intermediaries, each of which inherently decreases the reliability of the global system. A secure layer is important, but in order to have more than a secure layer on top of a fragile foundation layer it is necessary to design security and robustness into the system interactions. The Internet requires end-user protocol-based cooperation, and the Power Grid requires multiple control centers; the Stanford URI is working on architectures which require neither of these limitations. In particular, this URI program is developing sophisticated mathematical models and systematic tools, and using them not only for post-mortem analysis of attacks and failures but also to design and implement new protocols which are resistant to these modes of failure.

This work is undergoing transition to current communication networks, and will have significant application to future military mixed wired and wireless command and control networks. The focus is on attacks at the network infrastructure level, not on attacks on the computers at the edge of the network.

2 Congestion and Buffering in Wired networks

We have studied the problem of designing globally stable, scalable congestion control algorithms for the Internet. Prior work primarily used linear stability as the criterion for such a design. Global stability has been studied only for single node, single source problems. In our work, we have obtained conditions for a general topology network accessed by sources with heterogeneous delays. We obtain a sufficient condition for global stability in terms of the increase and decrease parameters of the congestion control algorithm and the price functions used at the links.

The key idea in our recent work is to first show that the source rates are both upper and lower bounded, and then use these bounds in Razumikhin's theorem to derive conditions for global stability. However, a stumbling block in extending earlier results to a general network is the difficulty in obtaining reasonable bounds on the source rates and in finding an appropriate Lyapunov-Razumikhin function. We take a significant step in this direction by finding a Lyapunov-Razumikhin function that provides global stability conditions for a general topology network with heterogeneous delays.

The global stability condition derived thereby is delay-independent, and is given in terms of the increase and decrease parameters and a parameter of the price function. When the condition holds, the network is globally stable for all values of fixed communication delays and controller gains. It is different from most prior works, where the conditions are given in term of the gains and the delays. Since our global stability condition is delay-independent, the network is robust to the delays and the gains used by users in the network. On the other hand, our stability condition restricts the possible choices for the utility functions and the price functions, whereas earlier stability conditions like work for general utility functions. Characterizing the stability region when our condition is violated, but the local stability condition still holds, still an open problem. Our simulation results indicate that

the region of attraction could be large under such a scenario.

We show that one can obtain conditions for global stability that relate the parameters of the congestion algorithm to the parameters of the price functions used at the links of the network. We further considered a two-phase algorithm, with a slow-start phase followed by a congestion-avoidance phase, as in today's version of TCP-Reno, and showed that a three-phase approximation of this two-phase algorithm is still globally, asymptotically stable under the same conditions on the congestion control parameters.

2.1 Sizing issues in buffer routers

Large buffers in Internet routers often limit achievable throughput, requiring the use of off-chip DRAMs. A standard guideline used for buffer size design is $B = RTT \times C$, where C is the capacity of the link and RTT is the round trip time. Recently, this design rule has been questioned and the use of small buffer routers validated based on statistical multiplexing effects.

However, these results have been based on *static* network simulations with *fixed number of flows*. In our work, we have completed far more extensive network simulations evaluating the accuracy of these results in a *dynamic* environment where file flows arrive and depart, i.e., flows numbers are not fixed. More specifically, we assess the performance of dynamic networks with very small buffers, with the end-user in mind. As flows arrive and depart, link utilization should not be considered as the most important factor in the design of the network. In a static network, where the number of flows is fixed, utilization and goodput have a direct correspondence as each user sees an average throughput of $\frac{C \times U}{N}$. In a dynamic network, the number of flows is time-varying: there is no such correspondence between the link utilization and the end-to-end throughput. Therefore, we directly calculate end-to-end throughput seen by the users and use this as a metric for evaluating performance.

For completeness we first completed simulations with fixed number of long flows. We then showed that in this case smaller buffers can indeed be used without any significant effects on throughput. We then further showed that Poisson pacing of TCP is not necessary. In fact, our simulations demonstrate that the effect of short flows and RTT variations create sufficient randomness to ensure high link utilization.

The current Internet consists of extremely fast core routers and slow edge routers. The edge routers switch packets at a rate which is several orders of magnitude smaller than the core routers. Our simulations show that in this case, very small buffers can be used without affecting the throughput, even in a dynamic network.

We have also considered the case where edge routers switch packets at a rate comparable to that of core routers, that is, there are no access bandwidth constraints. In this case, our simulation results demonstrate that if the network is moderately congested, then increasing the buffer size will in fact result in a substantial increase in overall throughput. As an example, when the load on a 100Mbps

link is approximately 80%, we find an increase in average throughput of between 60%-100% as the buffer size is increased from 20 to 1000 packets. We have now showed that TCP-pacing does not improve average throughput significantly.

Considering the same architecture as in the preceding discussion, we show that under mildly loaded conditions (load less than 50%), increased buffer sizes do not lead to increased throughput. That is, small buffers are adequate only when the core router is guaranteed to operate at 50% load, or less. In summary, in contrast to previous work, our simulations show that actual performance of routers with small buffers depends on the type of Internet architecture assumed.

2.2 Fluid model development and analysis of priority processing schemes in the Internet

Previous simulation studies have shown that providing a simple priority to short flows in Internet routers can dramatically reduce their mean delay while having little impact on the long flows that carry the bulk of the Internet traffic. We have proposed simple fluid models that can be used to quantify these observations. These fluid models are justified by showing that stochastic models of resource-sharing among TCP flows converge to these fluid models when the router capacity and the number of users are large.

We showed that a Shortest Remaining Processing Time (SRPT) scheme dramatically improves short flow performance, while having little impact on long flows. This scheme requires the router to estimate whether a flow is short or long, which is not feasible given that routers do not have access to per-flow information. Alternatively, using simple sampling techniques, it can be determined fairly accurately whether a flow is long or short. Assuming such a mechanism exists and can be easily implemented, we evaluate the performance of such priority processing schemes analytically, and further strengthen our conclusions via simulations.

Without priorities, the nature of bandwidth sharing in the Internet favors long flows. A more equitable sharing discipline can be approximated by discriminatory processor sharing (DPS). Stochastic analysis of DPS is extremely difficult and closed-form solutions exist only for exponentially distributed service times. However, in a system with a large number of files and a large server capacity, such as the Internet, some form of the law of large numbers can be applied and the resulting stochastic system can be approximated by a deterministic system that can be modeled by a set of differential equations. We have proposed such fluid flow models to capture the resource sharing character of TCP flows in the Internet. These models consider the impact of access bandwidth constraints. Using these models, we showed analytically that a stochastic model for DPS converges to the fluid limit in a large system. This further characterizes the speed of convergence of the fluid limit to equilibrium.

2.3 Connection-level stability analysis in the Internet

In this work, we have studied connection-level models of file transfer requests in the Internet, where connection arrivals to each route occur according to Poisson processes and the file-sizes have phase-type distributions. We use Sum-of-Squares techniques to construct Lyapunov functions satisfying Foster's condition for stochastic stability.

3 Scheduling and Resource Allocation in Wireless Networks

3.1 A Large Deviations Analysis of Scheduling

In [37] we consider a cellular network consisting of a base station and N receivers. The channel states of the receivers are assumed to be identical and independent of each other. The goal is to compare the throughput of two different scheduling policies (a queue-length-based policy and a greedy scheduling policy) given an upper bound on the queue overflow probability or the delay violation probability. We consider a multi-state channel model, where each channel is assumed to be in one of L states. Given an upper bound on the queue overflow probability or an upper bound on the delay violation probability, we show that the total network throughput of the queue-length-based policy is no less than the throughput of the greedy policy for all N . We also obtain a lower bound on the throughput of the queue-length-based policy. For sufficiently large N , the lower bound is shown to be tight, strictly increasing with N , and strictly larger than the throughput of the greedy policy. Further, for a simple multi-state channel model (on-off channel), we prove that the lower bound is tight for all N .

Multiuser wireless scheduling has received much attention in recent years. Consider a cellular network consisting of a base station and N users (receivers), where the base station maintains N separate queues, one corresponding to each user. Assume time is slotted and the channel states of the receivers at each time slot are known at the base station. Then, the base station can decide which queues to serve according to their channel states. We considered the case where the base station operates in a TDMA fashion, i.e., the base station can serve only one queue in each time slot. Two scheduling policies have been widely studied in the literature: (i) the base station serves the user with the best (weighted) channel state (opportunistic scheduling) [34, 19]; or (ii) serve the one with the best queue-length-weighted channel state (queue-length based (QLB) scheduling) [31, 11, 26, 27, 7, 4, 21]. While the QLB scheduling is throughput optimal (i.e., can stabilize any set of user throughputs that can be stabilized by any other algorithm), opportunistic scheduling maximizes the total network throughput if all queues are continuously backlogged. If the arrival rates to the users are identical and the channel state distributions to the receivers are identical, then these two scheduling policies have the same stability region.

While stability is the first concern of scheduling policies, quality-of-service (QoS) is equally

important in applications. For example, we may require the queue overflow probability to be small or require small delays. The performance of different scheduling policies under QoS constraints has received much attention recently. For reasons of analytical tractability, much of the prior work assumes that the channels to all the receivers are independent and statistically identical. Under this assumption, and assuming identical user utilities, opportunistic scheduling policies become greedy policies in which the base station transmits to the receiver with the best channel state. In [25], the author studies a simple network consisting of two users where the channels are assumed to be independent, identically distributed ON-OFF channels. Using large-deviations techniques, it is shown that the total network throughput of the QLB policy is larger than the throughput of the greedy policy under the queue overflow constraint. In [11], a wireless network with N users and ON-OFF channels is considered. It is assumed that the arrivals are identical and Poisson, and the capacity when the channel is ON is one packet per time slot. It is then shown that, when the number of users increases from N to $2N$, the expected sum of queue lengths is non-increasing under the QLB policy, while it increases linearly under the greedy policy. Further, in [8], the behavior of the greedy policy for Rayleigh fading channels is studied and it is shown that under a delay constraint, the total network throughput of the greedy policy increases initially with the number of users, but eventually decreases and goes to zero when the number of the users is sufficiently large.

Motivated by these prior results, in our work reported in [37], we study the performance of the two scheduling policies (greedy and QLB) for a wireless network with multi-state channels and constant arrivals. Using sample-path large-deviations techniques that have been used in [5], [25] and [29], we obtain the following results:

1. Assuming a multi-state channel model and a constant arrival rate in each time slot, under the QLB policy, we compute a lower bound on the large-deviations exponent of the probability that at least one queue in the network exceeds a large threshold. We obtain lower bounds on the maximum network throughput under the QoS constraints, and for large N , the lower bounds are tight, strictly increasing, and strictly greater than the throughput of the greedy policy. For the ON-OFF channel model, we prove that the lower bounds are tight for all N . It was conjectured that in [25] that, for the ON-OFF channel model, the complexity of the calculation of the large-deviations exponent increases exponentially with increasing N , but we show here that a simple closed-form expression can be obtained.
2. Consider ON-OFF channels and the QLB policy. In [11], under the assumption that the channel capacity is one packet per time slot, for a different model, it is shown the expected sum of the queue lengths is nondecreasing when the number of users increases from N to $2N$. For the ON-OFF channel model, we show that the maximum network throughput is strictly increasing in N under the delay-violation constraint or queue overflow constraint. Our result does not only compare performance with N users and $2N$ users, but at all intermediate values

as well. Our result also holds even when the capacity of the network is greater than one packet-per-slot. Further, for the general multi-state channel model, the maximum throughput is shown to be strictly increasing with N for large N .

3. For the greedy policy, we analytically show that the throughput goes to a constant under the queue overflow constraint, and decreases to zero under the delay violation constraint. This result holds for the general multi-state channel model, and is consistent with the numerical results for Rayleigh fading channels in [8].
4. Under the QoS constraints, we show that the throughput of the QLB scheduling policy is no less than the throughput of the greedy policy. This conclusion was also obtained in [25] for a two-user system and under the queue overflow constraint. Here, we prove that it is true for networks with N users ($N \geq 2$) and multi-state channels.

3.2 Distributed Fair Resource Allocation in Cellular Networks in the Presence of Heterogeneous Delays

In [38] consider the problem of allocating resources at a base station to many competing flows, where each flow is intended for a different receiver. The channel conditions may be time-varying and different for different receivers. It has been shown in [8] that in a delay-free network, a combination of queue-length-based scheduling at the base station and congestion control at the end users can guarantee queue-length stability and fair resource allocation. We extended this result to wireless networks where the congestion information from the base station is received with a feedback delay at the transmitters. The delays can be heterogenous (i.e., different transmitters may have different round-trip delays) and time-varying, but are assumed to be upper-bounded, with possibly very large upper bounds. We showed that the joint congestion control-scheduling algorithm continues to be stable and continues to provide a fair allocation of the network resources.

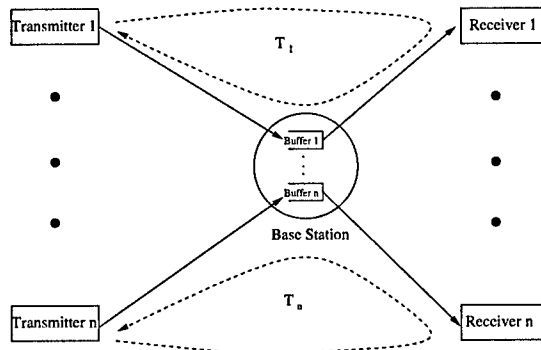


Figure 1: Network with feedback delays. The channel from the base station to the receivers is time-varying.

We have studied the problem of fair allocation of resources in the downlink of a cellular wireless network consisting of a single base station and many receivers (see Figure 1). The data destined for each receiver is maintained in a separate buffer. The arrivals to the buffers are determined via a congestion control mechanism. We assume that the time is slotted. The channels between the base station and the receivers are assumed to have random time-varying gains which are independent from one time-slot to the next. The independence assumption can be relaxed easily, but we use it here for ease of exposition. The goal is to allocate the network capacity fairly among the users, in accordance with the needs of the users, while exploiting the time-variations in the channel conditions. We associate a utility function with each user that is a concave, increasing function of the mean service that it receives from the network. In an earlier paper [8], it was shown that a combination of Internet-style congestion control at the end-users and queue-length based scheduling at the base station achieves the goal of fair and stabilizing resource allocation. This result is somewhat surprising since the resource constraints in the case of a wireless network are very different from the linear constraints in the case of the Internet [28]. The relative merits of congestion control-based resource allocation scheme as compared to other resource allocation schemes for cellular networks are discussed in [8]. Several other works in the same context are [30, 18, 20]. However, none of these works explicitly include the effect of feedback delay in their analysis. One of the reasons that delay is not important in these other works is that a specific scheduling algorithm is used in the network which allows the congestion control to be based only on the queue length at the entry node of each source. However, we considered a situation where such scheduling is not used and where the bottleneck is at the cellular network while the sources may be located far away from the base station. An example of such a situation is a file transfer from a remote host over the downlink of a cellular network. We aim to consider the effect of this essential parameter on the fairness and stability properties of the algorithm presented in [8].

In [8], it is assumed that there are no delays in the transmission of packets from an end-user (transmitter) to the base station and in the transmission of congestion information from the base station back to the end users. But if we consider the case where the end users are connected to the base station through the Internet, then delays exist in both directions: there is a propagation delay τ_i^f from the end user i to the base station — we call it the forward delay of the end user i , and a propagation delay τ_i^b from the base station to the end user i — we call it the backward delay. It is well-known that the presence of delays may affect the performance of the network. For example, Internet congestion controllers which are globally stable for the delay-free network may become unstable if the feedback delays are large [28]. In our problem, when delays exist, the information the end users obtain will be “outdated” information. So the congestion information the users obtain at time t does not reflect the queue status at the base station at time t . So it is interesting to study a wireless network with delays and ask whether the conclusions of [8] still hold for wireless networks

with heterogeneous delays. We answer this question by showing that for a network with uniformly-bounded delays, which are potentially heterogeneous and time-varying, the algorithm of [8] is stable and can be used to approximate weighted- m fair allocation arbitrarily closely. We emphasize that the results hold for networks with arbitrarily large, but bounded time-varying delays. So even if the end users can only get very old feedback information from the base station, the network is still stable and can eventually reach the fair resource allocation.

3.3 Simultaneous Routing and Resource Allocation

In wireless data networks the optimal routing of data depends on the link capacities which, in turn, are determined by the allocation of communications resources (such as transmit powers and signal bandwidths) to the links. The optimal performance of the network can only be achieved by simultaneous optimization of routing and resource allocation.

The paper [120] studies the simultaneous routing and resource allocation problem and exploits problem structure to derive efficient solution methods. We use a capacitated multicommodity flow model for the data flows in the network. We assume that the capacity of a wireless link is a concave and increasing function of the communications resources allocated to the link (TDMA and FDMA systems), and the communications resources for groups of links are limited. These assumptions allow us to formulate the simultaneous routing and resource allocation problem as a convex optimization problem over the network flow variables and the communications variables. These two sets of variables are coupled only through the link capacity constraints. We exploit this separable structure by dual decomposition. The resulting solution method attains the optimal coordination of data routing in the network layer and resource allocation in the radio control layer via pricing on the link capacities.

In [118], we generalize the simultaneous routing and resource allocation formulation to include CDMA wireless systems. Although link capacity constraints of CDMA systems are not jointly convex in rates and powers, we show that by using coordinate projections or transformations, the simultaneous routing and power allocation problem can still be formulated as (in systems with interference cancellation) or approximated by (in systems without interference cancellation) a convex optimization problem which can be solved very efficiently. We also propose a heuristic link-removal procedure based on the convex approximation to further improve the system performance.

4 Fault Diagnosis over Packet Dropping Networks

There are several challenges that arise when trying to perform management and control over unreliable, possibly heterogeneous networks. The major concern is the fact that, due to the nature of network links, observations may be delayed, lost or received out of order. In such case, diagnosers

or controllers that use this underlying network infrastructure need to be able to cope with the unreliability of the communication links in an effective and reliable manner. Within the context of this URI project, we have been exploring a number of directions that aim to achieve this ultimate goal.

4.1 Probabilistic Fault Detection and Identification

Another direction that we have been pursuing is along the lines of our work on systematic and efficient methodologies for fault management in dynamic systems. For example, we have developed probabilistic schemes for detecting permanent or transient functional changes (faults) in large-scale discrete event systems that can be modeled as finite-state machines. In one particular setup, the detector observes the frequencies with which states are occupied and detects faults by analyzing the deviation between the expected frequencies and the actual measurements. These features can be useful in distributed or networked settings where the input-state order may not be known and, at this point, we are considering applications of these ideas in the context of statistical methods for network security and intrusion detection. This work appeared as an invited paper during the 2002 Conference on Decision and Control; an extended version of it also appeared in the IEEE Transactions on Automatic Control.

More recently, we have begun the investigation of schemes for observing/diagnosing/controlling systems or networks under unreliable information that might arise due to permanent or transient faults in the system sensors. More specifically, we have developed a probabilistic methodology for failure diagnosis in finite state machines given a sequence of unreliable (possibly corrupted) observations. Assuming prior knowledge of the input probability distribution but no knowledge of the actual input sequence, the core problem we considered aimed at choosing from a pool of known, deterministic finite state machines (FSMs) the one that most likely matches a given (output) sequence of observations. The main challenge is that errors, such as symbol insertions, deletions, and transpositions, may corrupt the observed output sequence; the cause of these errors could be a faulty sensor or problems encountered in the communication channels or network links connecting the system sensors with the diagnoser/observer. Given the possibly erroneous output sequence of observations, we have proposed an efficient recursive algorithm for obtaining the most likely underlying FSM. We have illustrated the proposed methodology using as an example the diagnosis (identification) of a communication protocol.

Along these lines, we have also been able to make connections with the literature on hidden Markov models. To this end, we are currently trying to understand the role of reduced-order models and the role of modeling methodologies such as hidden Markov models or the influence model. We believe that this work will have important practical implications because it relates directly to the issue of sensor reliability and cost (i.e., the task of determining the required levels of reliability for the system sensors in order to guarantee a certain level of performance).

4.2 Distributed Symmetric Function Computation in Noisy Wireless Sensor Networks with Binary Data

With the wide availability of inexpensive wireless technology and sensing hardware, wireless sensor networks are expected to become commonplace because of their broad range of potential applications. A wireless sensor network consists of sensors that have sensing, computation and wireless communication capabilities. Each sensor monitors the environment surrounding it, collects and processes data, and when appropriate transmits information so as to cooperatively achieve a global detection objective. We have considered the common situation where there is a single fusion center, and the network goal is to cooperatively provide information to this fusion center so it can compute some function of the sensor measurements.

We have investigated this problem in multi-hop networks with noisy communication channels where the measurement of each sensor consists of one bit. We consider a sensor network consisting of n sensors, each having a recorded bit, the sensor's measurement, which has been set to either "0" or "1". The goal of the fusion center is to compute a symmetric function of these bits; i.e., a function that depends only on the number of sensors that have a "1". Specifically, distributed symmetric function computation with binary data, which is also called a counting problem, is as follows: each node is in either state "1" or "0", and the fusion center needs to decide, using information transmitted from the network, the number of sensors in state "1".

The sensors convey information to the fusion center in a multi-hop fashion to enable the function computation. The problem studied is to minimize the total transmission energy used by the network when computing this function, subject to the constraint that this computation is correct with high probability.

When nothing is known about the structure of the function to be computed, all bits must to be transmitted to the fusion center, and this is purely a routing problem when the channels are reliable. When the wireless channels are unreliable, the use of channel coding (see, for example, [9]) makes it possible to convey information in a point-to-point fashion with arbitrarily small amounts of error. However, the use of point-to-point error-correction coding without any in-network processing may result in high energy cost and delay. Our focus is computation of symmetric functions in a noisy wireless sensor network when total energy consumption *is* a major concern.

The algorithms considered are related to the algorithms for distributed computation over noisy networks, which are studied in [10, 23, 24, 22, 17], and references within. In both problems, the goal is to compute the value of some function based on the information of the nodes. Our work is closely related to parity computation and threshold detection in noisy radio networks studied in [10] and [17], respectively, where a broadcast network is assumed, in which all nodes can hear all transmissions, and each node has a "1" or a "0". The goal in [10, 17] was to investigate the minimum number of transmissions required to compute the parity or decide whether the number of nodes in

state “1” has exceeded the threshold value. Note that parity and threshold detection are special cases of counting, since both of these are determined if we know how many nodes have “1”.

While the problems considered in [10] and [17] are similar to our problem, a major difference is that in our model, each node may not be able to hear every other node in the network. The reason for this is that energy consumption can be an important consideration in wireless networks and it is well-known that it can be reduced significantly if the transmissions are carried out in a multi-hop fashion. This is a consequence of the well-known propagation model used to model wireless communication channels, whereby the energy required to transmit over a distance of r is proportional to r^α , where $\alpha \geq 2$ is a constant depending upon the environment. Thus, instead of each sensor sending its information to the fusion center directly, it is more efficient from an energy consumption point of view to send the information through relay nodes. It may be possible to reduce energy consumption even further by using some form of in-network data processing. This may have further benefits; for instance, if all the sensor measurements are to be transmitted from the sensors to the fusion center, then relay nodes closer to the fusion center would be depleted of their energy faster than nodes that are further away from the fusion center. Thus, in-network processing to reduce the number of transmissions could be beneficial for eliminating hotspots. Fundamentally, this is the distinction between multi-hop wireless networks used for communication and multi-hop wireless networks used for sensing. In multi-hop wireless communication networks, the protocols are designed so that they are not application-specific, and therefore the network can support a constantly evolving set of applications. Contrasting this, in multi-hop sensor networks, the architecture and protocols can be designed for each specific application, exploiting its structure, to reduce the energy usage within the network. This is the motivation for the recent works reported in [12] and [15]. In [12], the authors have designed a block coding scheme to compress the amount of information to be transmitted in a sensor network computing some functions. In [15], the authors investigate the optimal computation time and the minimum energy consumption required to compute the maximum of the sensor measurements. However, the in-network processing that we consider is different from the processing considered in [12] and [15], where the communication channels are assumed to be reliable, and the processing is to primarily exploit the spatial correlation [15] or the spatio-temporal correlations [12]. In our problem, processing is required not only to reduce the redundancy in the information to be conveyed in the fusion center, but also to introduce some redundancy to combat the effect of the noisy channels in the sensor network. Our results show that the additional redundancy required to combat channel errors does not significantly negate the benefits of in-network computation used to eliminate redundancy in the information, and the combination of in-network computation and channel coding could reduce the number of transmissions required in multi-hop networks to the same order as the number required in single-hop networks.

We use the routing protocol in [12] along with ideas from distributed parity computation in

noisy networks ([10]) to devise near energy-optimal algorithms for counting in sensor networks. A key difference between our work and the work in [10] is that, in the case of sensor networks, the fusion center does not communicate directly with each of the sensors. Thus, local computation is necessary before conveying some aggregate information in a multi-hop fashion to the fusion center. The local computation in our case is not a simple parity computation as in [10] but as we will see later, the network needs to compute the number of sensors in each local neighborhood (called a *cell*) that have seen a “1”. Further, we require that the computation be accurate uniformly over all cells. In addition, we will show that error-correction coding is required in the algorithms to minimize the energy required for counting.

We assume the wireless channels are binary symmetric channels with a probability of error p , and that each sensor uses r^α units of energy to transmit each bit, where r is the transmission range of the sensor. Using the above ideas, we first study the case where each sensor has only one observation to report, and show that the amount of energy required for counting (i.e., detecting the number of sensors seeing a “1”) is $O\left(n(\log \log n) \left(\sqrt{\frac{\log n}{n}}\right)^\alpha\right)$, where n is the number of sensors in the network. We also show that any algorithm satisfying the performance constraints must necessarily have energy usage $\Omega\left(n \left(\sqrt{\frac{\log n}{n}}\right)^\alpha\right)$. Then, we consider the case where the sensor network observes N events, and each node records one bit per event, thus having N bits to convey. The fusion center now wants to compute N symmetric functions, one for each of the events. We then extend to the case where each sensor has N binary observations, and the symmetric function needs to be computed for each observation. We show that the total transmission energy consumption can be reduced to $O\left(n \left(\max\left\{1, \frac{\log \log n}{N}\right\}\right) \left(\sqrt{\frac{\log n}{n}}\right)^\alpha\right)$ per observation. When $N = \Omega(\log \log n)$, the energy consumption is $\Theta\left(n \left(\sqrt{\frac{\log n}{n}}\right)^\alpha\right)$ per observation, which is a tight bound. If we only want to know roughly (a notion made precise in [39]) how many sensors have “1”. The answer can be obtained with the transmission energy consumption $\Theta\left(n \left(\sqrt{\frac{\log n}{n}}\right)^\alpha\right)$.

5 Decentralized Control

5.1 Control over Networks

The first line of research we have been pursuing relates to the study of the fundamental performance limitations of control methodologies that use existing network infrastructure as their communications backbone. For instance, by modeling a packet dropping network as an erasure channel and by focusing on bounded variance stabilization schemes, one can study the problem of plant stabilization despite message delays, packet drops, quantization noise and measurement noise. Our initial work on this problem appeared as an invited paper in the 2002 Conference on Decision and Control and focused on the case when the system to be stabilized is a discrete-time linear time-invariant system,

the communication links (between sensors, controller(s) and actuators) are part of a packet dropping network in which transmissions are independent, and the network packets are large enough so that the effect of measurement quantization can be modeled by an additive white noise process. This work, which has been referenced extensively by many researchers, has also been extended to settings where the system to be stabilized is a continuous-time system, in which case one additional parameter that needs to be chosen is the sampling rate at which the sample-data controller is operating. We have been able to determine ways to optimally choose this and other parameters of this control problem, and our results are currently under revision in the IEE Proceedings on Control Theory and Applications.

Related to the task described above is our study of the effects of roundoff noise on our ability to detect and identify transient faults that affect the operation of control systems. This roundoff noise could arise due to finite precision limitations of our controllers or due to quantization that takes place when sending information of the underlying communication network (e.g., the Internet). Our analysis has provided insight that allows us to handle roundoff noise via explicit bounds on the precision needed to guarantee the correct identification of the number of errors. Our analytical bounds can be very tight for certain choices of design parameters and can be used to provide guidance about the design of fault-tolerant systems.

More recently, our group has been focusing its efforts in extending these ideas to settings where the network delays between different packets are not independent. To this end, we have been trying to make connections with work on linear jump Markov systems. We are also interested in understanding how different network protocols (e.g., forward error correction or path diversity techniques) can be used to enable more effective controllers.

5.2 Decentralized Observation and Monitoring

Another direction that we have been pursuing within the context of this project relates to the construction of observers for switched systems under unknown or partially known inputs. This is a situation that arises frequently in practice as unknown inputs are used to represent uncertain system dynamics and faults or, in the case of decentralized systems, control signals generated by other controllers. Within this line of work, we have obtained methods for constructing reduced-order state observers for linear systems with unknown inputs. Apart from making connections with existing work on system invertibility and fault detection and identification, our approach provides a characterization of observers with delay, which eases the established necessary conditions for existence of unknown input observers with zero-delay. Our techniques are quite general in that they encompass the design of full-order observers via appropriate choices of design matrices.

Our work has also looked at challenges that arise in monitoring and controlling discrete event systems over unreliable networks. For instance, we have looked at decentralized failure diagnosis

schemes for systems that can be modeled as finite state machines. The specific scenario we considered consists of multiple local diagnosers, each with partial access to the outputs of the system under diagnosis. Our focus has been on designing a global coordinator which synchronizes with the local diagnosers at unspecified time intervals, and combines the local estimates in order to reach a final diagnosis. Under the assumption that the system and the local diagnosers are known, we have been able to analyze the effectiveness of simple types of global coordinators that operate without knowledge of the functionality of the system or the local diagnosers. In each case, we were able to derive conditions for finite-delay and zero-delay diagnosability, and to explore the trade-offs between the various schemes in terms of processing power and memory requirements on the global coordinator.

5.3 Decentralized Control

In order to develop systems that can coordinate with each other via communication the research in this program targets many new issues which are not present in the traditional feedback control scenario. These include distributed control decision based on local rather than global information, asynchronous information transmission, dynamic network topology, and scalability of algorithms to networks with large-numbers of nodes. Separate idealizations of the first two of these aspects of the problem have been possible in the past, although even these model problems lead to substantial difficulties in analysis. For example, one may use an idealized model of communication, and consider the simplest *decentralized control* problem. A general formulation of this problem can be reduced to one of structured control synthesis, a problem for which a systematic approach is lacking in the current literature.

Conversely, one may assume a centralized information pattern and consider the centralized control problem subject to *asynchronous communication*. In this case, instead of data consisting of continuous signals, data is now transmitted in packets. Furthermore, packets are subject to loss or delay, large packets may be fragmented and require reassembly, and packet streams may be received out of order. Control systems must be designed which are robust to these occurrences.

Given a particular plant and a constraint set of allowable decentralized controllers, one would like to determine if the associated control problem is, in a certain sense, easily solvable. The paper [121] develops a clear and precise characterization of when this is so. The notion of *quadratic invariance* is introduced in that paper, and is outlined here.

We suppose we have a linear plant G , and a subspace of admissible controllers S , which captures any sparsity constraints on the controller. The set S is called quadratically invariant if KGK is an element of S for all K in S . The paper shows that, if the constraint set has this property, then a controller which minimizes any norm of the closed-loop system may be efficiently found.

The area of decentralized control systems has been a source of challenging problems for many

years. Starting with work on team theory, there have been many results showing that problems with certain information structures may be solved, and recently many papers have developed specific optimization methods to address these problems. Examples include decentralized control where the systems are chained in a particular way, as well as arrays of systems where information satisfies certain delay requirements. Our work provides a unifying framework in which to analyze these systems. So far, all of the known solvable problems to which we have applied this theory have been found to be quadratically invariant.

There are many links here to other active areas of research within this project, in particular to the work on the decentralized congestion control mechanisms used by TCP in the Internet. Further work remains, as there are important questions of computational complexity, since many decentralized control problems are known to be intractable. It is also known that many decentralized linear control systems have optimal controllers which are nonlinear. Our work addresses some extremely fundamental issues which are a central concern, and produces results which are both theoretically important and practically relevant.

5.4 Monitoring and Control of Power System Dynamics

We have addressed monitoring and control of system-wide electromechanical (or ‘swing’) dynamics in power systems, as well as the dynamics of auction-based electricity markets. We showed that observer-based power system monitors can be used to estimate the full state of the system as well as identify and isolate a number of events (e.g., faults) using only sparse local measurements, all in the presence of various system disturbances. This work also develops and exploits a spatio-temporally integrated view of electromechanical dynamics. This contrasts with the traditional approach of either studying temporal variations at fixed spatial points or investigating spatial variations of specified (e.g., modal) temporal behavior. We use a continuum model of the swing dynamics to expose the wave-like propagation of electromechanical disturbances and to gain insight for the design of controls. This leads to strategies for decentralized control of these electromechanical waves, drawing on prototype controllers found in electromagnetic transmission line theory (e.g., matched-impedance terminations) and active vibration damping (e.g., energy-absorbing controllers and vibration isolators). Finally, we have proposed various controllers to realize quenching or confining-and-quenching strategies, and tested these in simulations of a 179-bus reduced-order representation of the power grid of the western US and Canada.

6 Complexity and robustness in complex networks

Recent progress in systems biology and network-based technological systems, together with new mathematical theories, has revealed generalized principles that shed new light on complex networks,

and confirmed the observations that an inherent feature of complex multiscale systems is that they are "robust yet fragile" (RYF). They are both intrinsically robust under most normative conditions and yet can be extremely sensitive to certain perturbations in their environment and component parts. This RYF feature provides a new paradigm for thinking about complexity and evolution across a broad range of phenomena and scales from computer networks to immune systems, from power grids and cancers to ecosystems, financial markets and human societies. While this research draws heavily on systems and control theory, most papers have appeared in biology, networking, and physics journals. In this section, we will review recent progress in theory and applications aimed at an engineering audience.

This research builds on insights about the fundamental nature of complex biological and technological networks that can now be drawn from the convergence of three research themes. 1) Molecular biology has provided a detailed description of much of the components of biological networks, and with the growing attention to systems biology the organizational principles of these networks are becoming increasingly apparent. 2) Advanced technology has provided engineering examples of networks with complexity approaching that of biology. While the components differ from biology, we have found striking convergence at the network level of architecture and the role of layering, protocols, and feedback control in structuring complex multiscale modularity. Our research is leading to new theories of the Internet and to new protocols that are being tested and deployed for high performance scientific computing. 3) Most importantly, there is a new mathematical framework for the study of complex networks that suggests that this apparent network-level evolutionary convergence both within biology and between biology and technology is not accidental, but follows necessarily from the requirements that both biology and technology be efficient, adaptive, evolvable, and robust to perturbations in their environment and component parts. This theory builds on and integrates decades of research in pure and applied mathematics with engineering, and specifically with robust control theory.

Through evolution and natural selection or by deliberate design, such systems exhibit highly functional and symbiotic interactions of extremely heterogeneous components, the very essence of "complexity". At the same time this resulting organization allows, and even facilitates, severe fragility to cascading failure triggered by relatively small perturbations. Thus robustness and fragility are deeply intertwined in biological systems, and in fact the mechanisms that create their extraordinary robustness are also responsible for their greatest fragilities. Our highly regulated and efficient metabolism evolved when life was physical challenging and food was often scarce. In a modern lifestyle, this robust metabolism can contribute to obesity and diabetes. More generally, our highly controlled physiology creates an ideal ecosystem for parasites, who hijack our robust cellular machinery for their own purposes. Our immune system prevents most infections but can cause devastating autoimmune diseases, including a type of diabetes. Our complex physiology requires robust develop-

ment and regenerative capacity in the adult, but this very robustness at the cellular level is turned against us in cancer. We protect ourselves in highly organized and complex societies which facilitate spread of epidemics and destruction of our ecosystems. We rely on ever advancing technologies, but these confer both benefits and horrors previously unimaginable. This universal "robust yet fragile" (RYF) nature of complex systems is well-known to experts such as physicians and systems engineers, but has been systematically studied in any unified way only recently. It is now clear that it must be treated explicitly in any theory that hopes to explain the emergence of biological complexity, and indeed is at the heart of complexity itself.

These RYF features appear on all time and space scales, from the tiniest microbes and cellular subsystems up to global ecosystems, and also -we believe- to human social systems, and from the oldest known history of the evolution of life through human evolution to our latest technological innovations. Typically, our networks protect us, which is a major reason for their existence. But in addition to cancer, epidemics, and chronic auto-immune disease, the rare but catastrophic market crashes, terrorist attacks, large power outages, computer network virus epidemics, and devastating fires, etc, remind us that our complexity always comes at a price. Statistics reveal that most dollars and lives lost in natural and technological disasters happen in just a small subset of the very largest events, while the typical event is so small as to usually go unreported. The emergence of complexity can be largely seen as a spiral of new challenges and opportunities which organisms exploit, but lead to new fragilities, often to novel perturbations. These are met with increasing complexity and robustness, which in turn creates new opportunities but also new fragilities, and so on. This is not an inexorable trend to greater complexity, however, as there are numerous examples of lineages evolving increasing simplicity in response to less uncertain environments. This is particularly true of parasites that rely on their hosts to control fluctuations in their microenvironment, thus shielding them from the larger perturbations that their hosts experience.

It is only fairly recently, and particularly the last few decades, that human technology has become focused not just on robustness, but on architectures that facilitate the evolution of new capabilities and the scaling to large system sizes. Protocol-based multilayer modular design is permeating advanced technologies of all kinds, but the Internet remains perhaps the most well-known example. It is also particularly suitable for our purposes for several reasons. The Internet, and cybertechnology generally, are unprecedented in the extent to which their features parallel biology. Their most salient features are often hidden from the user and thus as metaphors are often terribly misleading, yet are extremely useful when right. Only cybertechnology has the potential to rival biotechnology in accelerating the human/technology evolution, and the combined RYF spiral could have profound consequence. The most consistent, coherent, and salient features of all complex technologies are their protocols. To engineers, the term "protocol" is the set of rules by which components interact to create system-level functionality. Indeed, in advanced technologies, and

we believe in the organization of cells and organisms, the protocols are more fundamental than the modules whose interconnection they facilitate, although they often are obscured by the overwhelming details that now characterize experimental results in biology. A central feature of efficient, protocol-based systems is that, provided they obey the protocols, modules can be exchanged. The details are less important here than the consequences, which are the system-level robustness and evolvability that these protocols facilitate. New and even radically different hardware is easily incorporated at the lowest physical layers, and even more radically varying applications are enabled at the highest layers. Ironically as in biology, it is these transient elements of hardware and application software that are most visible to the user, while the far more fundamental and persistent infrastructure is the core protocols, which by design remain largely hidden from the user.

A protocol-based organization facilitates coordination and integration of function to create coherent and global adaptation to variations in their components and environments on a vast range of time scales despite implementation mechanisms that are largely decentralized and asynchronous. The parallels here between the Internet and biology are particularly striking. The TCP/IP protocol suite enables adaptation and control on time scale from the sub-microsecond changes in physical media, to the millisecond-to-second changes in traffic flow, to the daily fluctuations in user interactions, to evolving hardware and application modules over years and decades. The remarkable robustness to changing circumstances and evolution of Internet-related technology could only have come about as the result of a highly structured and organized suite of relatively invariant and universally-shared, well-engineered protocols.

Similarly, a protocol-based architecture in biology and its control mechanisms facilitate both robustness and evolvability, despite massive impinging pressures and variation in the environment. With the most obvious example involving the table of codons, biology's universally shared set of protocols are more fundamental and invariant than the modules whose control and evolution they facilitate. Allostery, a huge suite of post-translational modifications, and the rapid changes in location of macromolecular modules enable adaptive responses to environmental signals or alterations on rapid time scales. Translational and transcriptional control and regulation of alternative splicing and editing act on somewhat longer time scales. On still longer time scales within and across generations, the sequences of the DNA itself can change, not only through random mutation, but also through highly structured and evolved mechanisms that facilitate the generation of adaptive diversity. Furthermore, as biologists dig deeper past the superficiality of sequence data into the complexity of regulation, they unearth additional layers of control that are fundamentally similar to those in advanced technologies. There is seemingly no limit to the ingenuity that biology uses in creating additional layers of sophisticated control. Now familiar examples range from RNA editing and alternative splicing to transposons, mismatch repair, and repetitive sequences to the cutting and pasting in the "arms race" of the immune system versus spirochete and trypanosome coat proteins.

Perhaps the most familiar example of lateral gene transfer in bacteria is possible because bacteria have a shared set of protocols that have even been quite appropriately described by some as the "bacterial Internet." Bacteria can simply grab DNA encoding new genes from other bacteria and incorporate it into their genome, just like computer users can buy a new computer and plug it into home or office networks. This "plug and play" modularity works because there is a shared set of protocols that allow even novel genetic material to be functional in an entirely new cellular setting. Plug and play DNA mobility and expression is further facilitated by integrons and plasmids. Thus, for example, bacteria can acquire antibiotic resistance on time-scales that would be vanishingly improbable by point mutations, an example of how rapid evolution of complexity is possible by Darwinian mechanisms. Natural selection can favor the evolution of whole protocol suites, and their interactions, which in turn massively accelerate the acquisition and sharing of functional adaptive change. Thus evolvability itself can be seen as the robustness of lineages, rather than organisms, on long time scales and to possibly large changes in the environment, indeed ones that would be lethal to organisms if they occurred rapidly. An important insight is that robustness and evolvability are generally not in conflict, and both are the product of systematic and organized control mechanisms.

The framework being developed here is radical in both its methodology and philosophical implication. Methodologically it draws on mathematics that is often not well known outside expert circles and in many cases had not traditionally been thought of as "applied." The mathematics tells us that robustness and fragility have conserved quantities, and we believe these will ultimately be of as much importance to understanding biological complexity as energy and entropy were to understanding the steam engine and mitochondria. The above views of "organized complexity" motivated by biology and engineering contrast sharply with that of "emergent complexity" that is more popular within the physical sciences. Highly Optimized/Organized Tolerance/Tradeoffs (HOT) has aimed to explain the issues of organized complexity, but emphasizing models and concepts such as lattices, cellular automata, spin glasses, phase transitions, criticality, chaos, fractals, scale-free networks, self-organization, and so on, that have been the inspiration for the "emergent" perspective. A side benefit of this largely pedagogical effort is it has led to apparently novel insights into RYF aspects of longstanding mysteries in physics, from coherent structures in shear flow turbulence and coupled oscillators, to the ubiquity of power laws, to the nature of quantum entanglement, to the origin of dissipation. Finally, the underlying mathematics may offer new tools to explore other problems in physics where RYF features may play a role, particularly involving multiple scales and organized structures and phenomena.

7 Network Reliability

We have investigated the reliability of networks operated in a distributed manner to changes in the network. Our work is based on the use of network coding, which allows both distributed coding and distributed implementations of cost optimization for the subgraphs used for network coding. We have considered two main issues: 1) the reliability of networks to packet losses 2) the cost efficiency of distributed coding and optimization under changes of topology, cost and traffic.

Packet losses in networks result from a variety of causes, which include congestion, buffer overflows, and, in wireless networks, link outage due to fading. Thus a method to ensure reliable communication is necessary, and the prevailing approach is for the receiver to send requests for the retransmission of lost packets over some feedback channel. There are, however, a number of drawbacks to such an approach, which are evident most notably in high-loss environments and for multicast connections. In both instances, many requests for retransmissions are usually required, which place an unnecessary load on the network and which may overwhelm the source. In the latter instance, there is the additional problem that retransmitted packets are often only of use to a subset of the receivers and are therefore redundant to the remainder. An approach that overcomes these drawbacks is to use erasure-correcting codes. Under such an approach, the original packets are reconstructed from those that are received and little or no feedback is required. This approach has been recently exemplified by digital fountain codes, which are fast, near-optimal erasure codes. Such codes can approach the capacity of connections over lossy packet networks, provided that the connection as a whole is viewed as a single channel and coding is performed only at the source node. But in lossy packet networks where all nodes have the capability for coding, such as overlay networks using UDP and wireless networks, there is no compelling reason to adopt this view, and a greater capacity can in fact be achieved if we do not.

We have developed a capacity-approaching coding scheme for unicast or multicast over lossy packet networks. In the scheme, all nodes perform coding, but do not wait for a full block of packets to be received before sending out coded packets. Rather, whenever they have a transmission opportunity, they form coded packets with random linear combinations of previously received packets. All coding and decoding operations in the scheme have polynomial complexity. Our analysis of the scheme has shown that it is not only capacity-approaching, but that the propagation of packets carrying innovative information follows that of a queueing network where every node acts as a stable MM1 queue. We are able consider networks with both lossy point-to-point and broadcast links, allowing us to model both wireline and wireless packet networks.

In the area of distributed optimization, we have presented decentralized algorithms that compute minimum-cost subgraphs for establishing multicast connections in networks that use coding. These algorithms, coupled with existing decentralized schemes for constructing network codes, constitute a fully decentralized approach for achieving minimum cost multicast. Our approach is in sharp

contrast to the prevailing approach based on approximation algorithms for the directed Steiner tree problem, which is suboptimal and generally assumes centralized computation with full network knowledge. We also have developed extensions beyond the basic problem of fixed-rate multicast in networks with directed point-to-point links, and consider the case of elastic rate demand as well as the problem of minimum energy multicast in wireless networks. For the case of optimization under changing conditions, we have given a formulation of the dynamic multicast problem for coded networks that lies within the framework of dynamic programming. Our formulation addresses the desired objective of finding minimum-cost time-varying subgraphs that can deliver continuous service to dynamic multicast groups in coded networks and, because it lies within the framework of dynamic programming, can be approached using methods developed for general dynamic programming problems. The solution that we propose uses such methods to approximate the optimal cost function, which is used to modify the objective function of an optimization that determines the multicast subgraph to use during each time interval. Depending upon the approximation that is used for the optimal cost function, this optimization conducted every time interval may be tractable and may even be amenable to decentralized computation.

7.1 Graph Structure and Similarity

We have investigated measures of graph similarity, developed a new measure, and applied it to the matching of graph fragments to their original locations in a parent graph. Measures of graph similarity have a broad array of applications, including comparing chemical structures, navigating complex networks like the World Wide Web, and more recently, analyzing different kinds of biological data. The research focuses on an interesting class of iterative algorithms that use the structural similarity of local neighborhoods to derive pairwise similarity scores between graph elements. Our new similarity measure uses a linear update to generate both node-node and edge-edge similarity scores, and has desirable convergence properties. The research also explores the application of our similarity measure to graph matching. We attempt to correctly position a subgraph within a parent graph using a maximum-weight matching algorithm applied to the similarity scores between the two graphs. Significant performance improvements are observed when the ‘topological’ information provided by the similarity measure is combined with additional information — such as partial labeling — about the attributes of the graph elements and their local neighborhoods. Further work is needed to explore various extensions of these ideas, including to the case of dynamically evolving graphs.

In other work, we study synchronization of complex random networks of nonlinear oscillators, with identical oscillators at the nodes interacting through ‘diffusive’ coupling across edges of the interconnection graph. Our random network is constructed by a generalized Erdos-Renyi method, so as to have specifiable expected-degree distribution. We present a sufficient condition for synchronization and a sufficient condition for desynchronization, stated in terms of the coupling strength

and the extreme values of the distribution of nontrivial eigenvalues of the graph Laplacian. We then determine the Laplacian eigenvalue distribution for the case of large random graphs through computation of the moments of the eigenvalue density function. The analysis is illustrated using a random network with a power-law expected-degree distribution and chaotic dynamics at each node. The mathematical structure of our problem is closely related to that of consensus problems in networks of agents, as well as the task of analyzing flocking/swarming conditions in a group of autonomous agents.

8 Impact of research and transitions

1) The first important algorithm to be transitioned is Approximate Fair Dropping (AFD). This is a new randomized algorithm that partitions the bandwidth of a link among the flows traversing the link. It builds on the CHOKe (choose and keep or choose and kill) algorithm, which is a simple algorithm for protecting TCP flows from UDP flows and enables the detection of flows which attempt to take up a disproportionate share of resources. The randomized nature of the algorithms not only make them simpler to implement, but also prevents users from predicting and attempting to spoof their behavior. The AFD algorithm is in discussion for implementation in the CISCO GSR12000 series of core routers.

2) Secondly, the FAST (Fast AQM, Scalable TCP) protects the TCP protocol from instabilities which current occur at high link speeds. These instabilities cause network throughput to drop to an extremely low level, and affect fast networks dramatically. Building on research from both the controls and networking community has led to this new protocol, which is both provably robust and scalable as well as incrementally deployable. In an experiment in November 2002, a speed of 8,609 megabits per second (Mbps) was achieved by using 10 simultaneous flows of data over routed paths, which is the largest aggregate throughput ever accomplished in such a configuration. FAST has been developed in significant part by Caltech, and has been transitioned through the Stanford Linear Accelerator Center (SLAC), working in partnership with the European Organization for Nuclear Research (CERN), and DataTAG, StarLight, TeraGrid, Cisco, and Level3.

3) This program has developed Distributed Random Coding (DRC), which combines the benefits of coding and routing into a single protocol. Here, instead of simply forwarding packets, nodes construct and forward algebraic combinations of inputs. This results in a network which has both significantly increased throughput as well as making it impossible for an observer to decode data transmitted by simply observing the network at a single point. This protocol has been implemented in a large-scale network testbed by Microsoft, working with Sprint.

All of the above protocols are backed up by theoretical analysis, with, for example, associated proofs of stability and convergence. The program has developed several further protocols for net-

work security, including mechanisms for covert message transmission via timing channels, protection against SYN flooding, the SIFT algorithm for prioritization in caches and buffers, and the SHRiNK method for monitoring extremely large-scale networks.

References

- [1] R. Pan, B. Prabhakar, K. Psounis: "CHOCe - A stateless active queue management scheme for approximating fair bandwidth allocation". *INFOCOM 2000*.
- [2] R. Pan, L. Breslau, B. Prabhakar, S. Shenker: "Approximate fairness through differential dropping," *ACM Computer Communications Review*, January 2002.
- [3] R. Pan, L. Breslau, B. Prabhakar, S. Shenker: "Approximate fair allocation of link bandwidth," *IEEE Micro*, January 2003.
- [4] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, and P. Whiting. CDMA data QoS scheduling on the forward link with variable channel conditions. Bell Laboratories Tech. Rep., April 2000.
- [5] D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis. Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach. In *IEEE Transactions on Automatic Control*, 43:315-335, March 1998.
- [6] S. Deb and R. Srikant. Global stability of congestion controllers for the internet. *IEEE Transactions on Automatic Control*, 48(6):1055-1060, June 2003.
- [7] A. Eryilmaz and R. Srikant, and J. Perkins. Stable scheduling policies for fading wireless channels. In *IEEE/ACM Transactions on Networking*, April 2005, pp. 411-424.
- [8] A. Eryilmaz and R. Srikant. Scheduling with Quality of Service Constraints over Rayleigh Fading Channels In *Proceedings of IEEE Conference on Decision and Control*, Dec. 2004.
- [9] R. G. Gallager. Information Theory and Reliable Communication. John Wiley & Sons, New York, 1968.
- [10] R. G. Gallager. Finding parity in a simple broadcast network. In *IEEE Transactions on Information Theory*, vol. 34, pp 176-180, 1988.
- [11] A. Ganti, E. Modiano and J. Tsitsiklis. Optimal Transmission Scheduling in Symmetric Communication Models with Intermittent Connectivity, 2004 Preprint.
- [12] A. Giridhar and P. R. Kumar. Computing and communicating functions over sensor networks. In *IEEE Journal on Selected Areas in Communications*, pp. 755-764, vol. 23, no. 4, April 2005.
- [13] P. Gupta and P. Kumar. Critical power for asymptotic connectivity in wireless network. In *Stochastic Analysis, Control, Optimization and Applications: a Volume in Honor of W.H.Fleming*, W. McEneaney, G. Yin and Q. Zhang, Eds., 1998
- [14] P. Gupta and P. Kumar. The capacity of wireless networks. In *IEEE transactions of Information Theory*, vol. 46, no.2, pp. 388-404, 2000.
- [15] N. Khude, A. Kumar and A. Karnik. Time and Energy Complexity of Distributed Computation in Wireless Sensor Networks. In *Proceedings of the IEEE Infocom*, 2005.

- [16] S. R. Kulkarni and P. Viswanath. A Deterministic Approach to Throughput Scaling in Wireless Networks. In *IEEE Trans. on Information Theory*, Vol. 50, No.6, pp. 1041-1049, June 2004.
- [17] E. Kushilevitz and Y. Mansour. Computation in Noisy Radio Networks In *Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*, pp. 236-243, 1998.
- [18] X. Lin and N. Shroff. The impact of imperfect scheduling on cross-layer rate control in wireless networks. In *Proceedings of IEEE Infocom 2005*, Miami, Florida, March 2005.
- [19] X. Liu, E. Chong, and N. Shroff. Opportunistic transmission scheduling with resource-sharing constraints in wireless networks. In *IEEE Journal on Selected Areas in Communications*, 19(10):2053 – 2064, October 2001.
- [20] M. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. In *Proceedings of IEEE Infocom 2005*, Miami, Florida, March 2005.
- [21] M. J. Neely, E. Modiano, and C. E. Rohrs. Power and server allocation in a multi-beam satellite with time varying channels. In *Proceedings of IEEE Infocom*, New York, NY. June 2002
- [22] S. Rajagopalan and L. J. Schulman. A Coding Theorem for Distributed Computation. In Proc. 26th STOC 790-799, 1994.
- [23] L. J. Schulman. Communication on Noisy Channels: A Coding Theorem for Computation. In *Proceeding of 33rd FOCS*, pp. 724-733, 1992.
- [24] L. J. Schulman. Deterministic Coding for Interactive Communication. In *Proceeding of the 25th Annual Symposium on Theory of Computing*, pp. 747-756, 1993.
- [25] S. Shakkottai. Effective Capacity and QoS for Wireless Scheduling, 2004 Preprint.
- [26] S. Shakkottai, R. Srikant, and A. Stolyar. Pathwise optimality of the exponential scheduling rule for wireless channels. In *Advances in Applied Probability*, 1021-1045, Dec. 2004.
- [27] S. Shakkottai and A. Stolyar. Scheduling for multiple flows sharing a time-varying channel: The exponential rule. In *Translations of the American Mathematical Society*, Series 2, A volume in memory of F. Karpelevich, Yu. M. Suhov, Editor, Vol. 207, 2002.
- [28] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [29] A. Stolyar and K. Ramanan. Largest weighted delay first scheduling: Large deviations and optimality. In *Ann. Appl. Probab.* 11:1-48, 2001.
- [30] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. Submitted.
- [31] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. In *IEEE Transactions on Information Theory*, 39:466-478, March 1993.
- [32] S. Toumpis and A. J. Goldsmith. Large wireless network under fading, mobility, and delay constraints. In *Proceedings of IEEE INFOCOM*, 2004.
- [33] G. Vinnicombe. On the stability of networks operating TCP-like congestion control. In *Proceedings of the IFAC World Congress*, Barcelona, Spain, 2002. University of Cambridge Technical Report CUED/F-INFENG/TR.398. Available at <http://www.eng.cam.ac.uk/~gv>.
- [34] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. In *IEEE Transactions on Information Theory*, 48(6):1277C1294, June 2002.

- [35] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, pp. 169–181, vol.10, no. 2, March 2004.
- [36] Ying, L., G. E. Dullerud, and R. Srikant, “Global Stability of Internet Congestion Controllers with Heterogeneous Delays,” *IEEE Transactions on Networking*; to appear June 2006.
- [37] Ying, L. , R. Srikant, A. Eryilmaz, and G. E. Dullerud, “A Large Deviations Analysis of Scheduling in Wireless Networks”, submitted to *IEEE Transactions on Information Theory*; conference version in *Proceedings of the IEEE Conference on Decision and Control*, 2005.
- [38] Ying, L. , R. Srikant, A. Eryilmaz, and G. E. Dullerud, “Distributed Fair Resource Allocation in Cellular Networks in the Presence of Heterogeneous Delays”, submitted to *IEEE Transactions on Automatic Control*; conference version in *Proceedings of the Intl. Symposium on Modelling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2005.
- [39] Ying, L. , R. Srikant, and G. E. Dullerud, “Distributed Symmetric Function Computation in Noisy Wireless Sensor Networks with Binary Data”, *Proceedings of the Intl. Symposium on Modelling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2006.
- [40] C. N. Hadjicostis, “Coding Approaches to Fault Tolerance in Combinational and Dynamic Systems.” Kluwer Academic Publishers, 2002.
- [41] S. Sundaram and C. N. Hadjicostis, “Non-Concurrent Error Detection and Correction in Switched Linear Controllers.” *International Workshop on Hybrid Systems: Computation and Control* (Series Lecture Notes in Computer Science, vol. 2993), pp. 585-599, Springer-Verlag, 2004.
- [42] S. Sundaram and C. N. Hadjicostis, “Error Detection and Correction in Switched Linear Controllers via Periodic and Non-Concurrent Checks.” To appear in *Automatica*.
- [43] S. Sundaram and C. N. Hadjicostis, “Comments on “Time-Delayed State Estimator for Linear Systems with Unknown Inputs”,” *International Journal on Control, Automation and Systems*, vol. 3, no. 4, pp. 646-647, December 2005.
- [44] C. N. Hadjicostis, “Probabilistic Fault Detection in Finite-State Machines Based on State Occupancy Measurements,” *IEEE Transactions on Automatic Control*, vol. 50, no. 12, pp. 2078-2083, December 2005.
- [45] Y. Wu and C. N. Hadjicostis, “Algebraic Approaches for Centralized and Distributed Fault Identification in Discrete Event Systems,” *IEEE Transactions on Automatic Control*, vol. 50, no. 12, pp. 2048-2053, December 2005.
- [46] C. N. Hadjicostis, “Aliasing Probability Calculations for Arbitrary Compaction under Independently Selected Random Test Vectors,” *IEEE Transactions on Computers*, vol. 54, no. 12, pp. 1614-1627, December 2005.
- [47] E. Athanasopoulou and C. N. Hadjicostis, “Probabilistic Approaches to Fault Detection in Networked Discrete Event Systems,” *IEEE Transactions on Neural Networks*, vol. 16, no. 5, pp. 1042-1053, September 2005 (Special Issue on Adaptive Learning Systems in Communication Networks).
- [48] C. N. Hadjicostis, “Finite-State Machine Embeddings for Non-Concurrent Error Detection and Identification,” *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 142-153, February 2005.

- [49] C. N. Hadjicostis and G. C. Verghese, "Coding Approaches to Fault Tolerance in Linear Dynamic Systems," *IEEE Transactions on Information Theory*, vol. 51, no. 1, pp. 210-228, January 2005.
- [50] C. N. Hadjicostis, "Non-Concurrent Error Detection and Identification in One-Hot Encoded FSMs," *Automatica*, vol. 40, pp. 1665-1676, 2004.
- [51] C. N. Hadjicostis, "Coding Techniques for Fault-Tolerant Parallel Prefix Computations in Abelian Groups," *The Computer Journal*, vol. 47, no. 3, pp. 329-341, 2004.
- [52] C. N. Hadjicostis, "Non-Concurrent Error Detection and Correction in Fault-Tolerant Linear Finite-State Machines," *IEEE Transactions on Automatic Control*, vol. 48, no. 12, pp. 2133-2140, December 2003.
- [53] C. N. Hadjicostis, "Non-Concurrent Error Detection and Correction in Fault-Tolerant Discrete-Time LTI Dynamic Systems," *IEEE Transactions on Circuits and Systems (I)*, vol. 50, no. 1, pp. 45-55, January 2003.
- [54] C. N. Hadjicostis and G. C. Verghese, "Encoded Dynamics for Fault Tolerance in Linear Finite-State Machines," *IEEE Transactions on Automatic Control*, vol. 47, no. 1, pp. 189-192, January 2002.
- [55] R. Touri, P. G. Voulgaris and C. N. Hadjicostis, "Time-Varying Power-Limited Preprocessing for Perfect Reconstruction of Binary Signals." To appear in the Proceedings of ACC 2006, the 2006 American Control Conference.
- [56] E. Athanasopoulou and C. N. Hadjicostis, "Decentralized Failure Diagnosis in Discrete Event Systems." To appear Proceedings of ACC 2006, the 2006 American Control Conference.
- [57] P. G. Voulgaris, C. N. Hadjicostis and R. Touri, "Encoder-Decoder Design for Perfect Reconstruction: A Robust Control Perspective." Proceedings of CDC/ECC 2005, the joint 44th IEEE Conference on Decision and Control and the European Control Conference 2005 (invited), pp. 2536-2541, Seville, Spain, 2005.
- [58] S. Sundaram and C. N. Hadjicostis, "On Delayed Observers for Linear Systems with Unknown Inputs." Proceedings of CDC/ECC 2005, the joint 44th IEEE Conference on Decision and Control and the European Control Conference 2005, pp. 7210-7215, Seville, Spain, 2005.
- [59] Athanasopoulou and C. N. Hadjicostis "Maximum Likelihood Diagnosis in Partially Observable Finite-State Machines." Proceedings of MED 2005, the 13th IEEE Mediterranean Conference on Control and Automation (invited), Limassol, Cyprus, 2005.
- [60] G. Takos and C. N. Hadjicostis, "Hierarchical Decentralized Fusion from Correlated Sensor Measurements." Proceedings of ICNSC 2005, the 2005 IEEE International Conference on Networking, Sensing and Control.
- [61] E. Athanasopoulou and C. N. Hadjicostis "Synchronization-Based Fault Detection in Discrete Event Systems." Proceedings of CDC 2004, the 43rd IEEE Conference on Decision and Control, Atlantis, Paradise Island, Bahamas, 2004.
- [62] Lingxi Li, C. N. Hadjicostis and R. Sreenivas, "Fault Detection and Identification in Petri Net Controllers." Proceedings of CDC 2004, the 43rd IEEE Conference on Decision and Control, Atlantis, Paradise Island, Bahamas, 2004.
- [63] C. N. Hadjicostis, "Finite-State Machine Embeddings for Non-Concurrent Error Error Detection and Identification," Proceedings of CDC 2003, the 42nd IEEE Conference on Decision and Control, vol. 4, pp. 3215-3220, Maui, Hawaii, 2003.

- [64] P. G. Voulgaris, C. N. Hadjicostis and R. Touri, "A Perfect Reconstruction Paradigm for Digital Communication," Proceedings of CDC 2003, the 42nd IEEE Conference on Decision and Control, vol. 4, pp. 4196-4201, Maui, Hawaii, 2003.
- [65] Y. Wu and C. N. Hadjicostis, "Distributed Non-Concurrent Fault Identification in Discrete Event Systems." Proceedings of CESA 2003, the Multiconference on Computational Engineering in Systems Applications (invited), Lille, France, 2003.
- [66] E. Athanasopoulou and C. N. Hadjicostis, "Aliasing Probability Calculations in Testing Sequential Circuits." Proceedings of MED 2003, the 11th IEEE Mediterranean Conference on Control and Automation, Rhodes, Greece, 2003.
- [67] C. N. Hadjicostis, "Encoded Finite-State Machines for Non-Concurrent Error Detection and Identification." Proceedings of ISCAS 2003, the 2003 IEEE International Symposium on Circuits and Systems (invited), vol. 3, pp. 858-861, Bangkok, Thailand, 2003.
- [68] C. N. Hadjicostis, "Aliasing Probability Calculations in Nonlinear Compactors." Proceedings of ISCAS 2003, the 2003 IEEE International Symposium on Circuits and Systems, vol. 5, pp. 529-532, Bangkok, Thailand, 2003.
- [69] C. N. Hadjicostis and R. Touri, "Feedback Control utilizing Packet Dropping Network Links." Proceedings of CDC 2002, the 41st IEEE Conference on Decision and Control (invited), vol. 2, pp. 1205-1210, Las Vegas, NV, 2002.
- [70] Y. Wu and C. N. Hadjicostis, "Non-Concurrent Fault Identification in Discrete Event Systems using Encoded Petri Net States." Proceedings of CDC 2002, the 41st IEEE Conference on Decision and Control (invited), vol. 4, pp. 4018-4023, Las Vegas, Nevada, 2002.
- [71] C. N. Hadjicostis, "Probabilistic Fault Detection in Finite-State Machines Based on State Occupancy Measurements." Proceedings of CDC 2002, the 41st IEEE Conference on Decision and Control (invited), vol. 4, pp. 3994-3999, Las Vegas, Nevada, 2002.
- [72] Papachristodoulou, A; Li, L; Doyle, JC, 2004. Methodological frameworks for large-scale network analysis and design, Computer Communication Review, Jul 2004
- [73] Li, L., D. Alderson, J. Doyle, and W. Willinger. 2004. A First-Principles Approach to Understanding the Internet's Router-level Topology. Proc. ACM Sigcomm, Computer Communication Review, Oct. 2004.
- [74] Csete M.E. and J.C. Doyle, 2004, Bow ties, metabolism, and disease, Trends in Biotechnology, Vol 22, Issue 9, pg. 446-450
- [75] J. Stelling, U. Sauer, Z. Szallasi, F. J. Doyle III, and J. Doyle, 2004, Robustness of cellular functions, Cell, October, 2004.
- [76] Willinger, W, D. Alderson, J.C. Doyle, and L.Li., 2004. More "Normal" Than Normal: Scaling Distributions and Complex Systems. Proceedings of the 2004 Winter Simulation Conference.
- [77] Shapiro, BE; Hucka, M; Finney, A; Doyle, J. 2004. MathSBML: a package for manipulating SBML-based biological models, Bioinformatics, Nov 1, 2004
- [78] Doyle, J; Csete, M, 2004. Imitation of life: How biology is inspiring computing, Nature. Oct 21 2004
- [79] Kitano, H; Oda, K; Kimura, T; Matsuoka, Y; Csete, M; Doyle, J; Muramatsu, M, 2004 .Metabolic syndrome and robustness tradeoffs, DIABETES 53: S6-S15 Suppl. 3, Dec 2004

- [80] H. El-Samad, H. Kurata, J.C. Doyle, C.A. Gross, and M. Khammash, 2005, Surviving Heat Shock: Control Strategies for Robustness and Performance, PNAS (8): 2736-2741 Feb 22, 2005
- [81] Jin C, Wei D, Low SH, Bunn J, Choe HD, Doyle JC, et al, FAST TCP: From theory to experiments IEEE Network 19 (1): 4-11 Jan-Feb 2005
- [82] Paganini F, Wang ZK, Doyle JC, et al. Congestion control for high performance, stability, and fairness in general networks, IEEE-ACM Transactions On Networking 13 (1): 43-56 Feb 2005
- [83] W. Willinger, J Doyle, Robustness and the Internet: Design and evolution, in Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies (Santa Fe Institute Studies on the Sciences of Complexity), Erica Jen, Editor
- [84] L. Chen, S. H. Low and J. C. Doyle. Joint congestion control and media access control design for wireless ad hoc networks, IEEE Infocom, Miami, FL, March 2005.
- [85] Wang JT, Li L, Low SH, Doyle JC. (2005) Cross-layer optimization in TCP/IP networks, IEEE-ACM Transactions On Networking 13 (3): 582-595 Jun 2005
- [86] Manning M, Carlson JM, Doyle J (2005) Highly optimized tolerance and power laws in dense and sparse resource regimes Physical Review E 72 (1): Art. No. 016108 Part 2 Jul 2005
- [87] R. Tanaka, T-M Yi, and J. Doyle (2005) Some protein interaction data do not exhibit power law statistics, FEBS letters, 579 (23): 5140-5144 Sep 26 2005
- [88] Doyle et al, (2005), The "Robust Yet Fragile" Nature of the Internet, P Natl Acad Sci USA. vol. 102 no. 41, October 11, 2005
- [89] Zhou T, Carlson JM, Doyle J (2005) Evolutionary dynamics and highly optimized tolerance, Journal Of Theoretical Biology 236 (4): 438-447 Oct 21 2005
- [90] L Li, D Alderson, JC Doyle, W Willinger (2005) Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications, Internet Math, to appear
- [91] D Alderson, L Li, W Willinger, JC Doyle (2005) Understanding Internet Topology: Principles, Models, and Validation, IEEE/ACM Transactions On Networking, Dec. 2005.
- [92] MA Moritz, ME Morais, LA Summerell, JM Carlson, J Doyle (2005) Wildfires, complexity, and highly optimized tolerance, P Natl Acad Sci USA, December 13, 2005; 102 (50),
- [93] T Brookings, JM Carlson, and J Doyle (2005) Three mechanisms for power laws on the Cayley tree, Phys. Rev. E 72, 056120
- [94] S. Boyd, P. Diaconis, P. Parrilo, and L. Xiao. Symmetry analysis of reversible Markov chains. To appear in *Internet Mathematics*, 2005.
- [95] S. Boyd, P. Diaconis, J. Sun, and L. Xiao. Fastest mixing Markov chain on a path. To appear in *The American Mathematical Monthly*, 2005.
- [96] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667-689, 2004.
- [97] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis and applications. Presented at *INFOCOM 2005*.
- [98] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Mixing times for random walks on geometric random graphs. Presented at *SIAM ANALCO 2005*.

- [99] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Analysis and optimization of randomized gossip algorithms. In *Proceedings of IEEE Conference on Decision and Control (CDC)*, volume 5, pages 5310–5315, Bahamas, December 2004.
- [100] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. Submitted to *IEEE Transactions on Information Theory*, March 2005.
- [101] C. Langbort, L. Xiao, R. D’Andrea, and S. Boyd. A decomposition approach to distributed analysis of networked systems. In *Proceedings of IEEE Conference on Decision and Control (CDC)*, volume 4, pages 3980–3985, Bahamas, December 2004.
- [102] J. Sun, S. Boyd, L. Xiao, and P. Diaconis. The fastest mixing Markov process on a graph and a connection to a maximum variance unfolding problem.
- [103] L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53(1):65–78, 2004.
- [104] L. Xiao and S. Boyd. Optimal scaling of a gradient method for distributed resource allocation. Revised for publication in *Journal of Optimization Theory and Applications*, February 2005. Available at http://www.stanford.edu/~boyd/fast_redstb.html.
- [105] L. Xiao, S. Boyd, and S.-J. Kim. Distributed average consensus with least-mean-square deviation. Submitted to *Journal of Parallel and Distributed Computing*, May 2005. Available at http://www.stanford.edu/~boyd/lms_consensus.html.
- [106] L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks (IPSN)*, pages 63–70, Los Angeles, CA, April 2005.
- [107] D. S. Lun, M. Medard, and R. Koetter. Efficient operation of wireless packet networks using network coding.. In Proc. International Workshop on Convergent Technologies (IWCT) 2005, June 2005. Invited paper.
- [108] S. Deb, M. Effros, T. Ho, D. R. Karger, R. Koetter, D. S. Lun, M. Medard, and N. Ratnakar. Network coding for wireless applications: A brief tutorial. In Proc. International Workshop on Wireless Ad-hoc Networks (IWWAN) 2005, May 2005. Invited paper.
- [109] D. S. Lun, N. Ratnakar, R. Koetter, M. Medard, E. Ahmed, and H. Lee. Achieving minimum cost multicast: A decentralized approach based on network coding. In Proc. IEEE Infocom 2005, March 2005.
- [110] D. S. Lun, M. Medard, and M. Effros. On coding for reliable communication over packet networks. In Proc. 42nd Annual Allerton Conference on Communication, Control, and Computing, September-October 2004. Invited paper.
- [111] Laura Zager, Graph Similarity and Matching, Masters thesis, EECS Department, MIT, May 2005.
- [112] Victor Preciado and George Verghese, "Synchronization in Generalized Erdos-Renyi Networks of Nonlinear Oscillators," IEEE Conf. on Decision and Control, Sevilla, Spain, December 2005.
- [113] Ernst Scholtz, Observer-Based Monitors and Distributed Wave Controllers for Electromechanical Disturbances in Power Systems, PhD thesis, EECS Department, MIT, August 2004.
- [114] Ernst Scholtz, Bernard Lesieutre and George Verghese, "Decentralized Controllers for Electromechanical Waves in Power Networks," Proc. 36th North American Power Symposium, August 2004, Moscow, Idaho.

- [115] Ernst Scholtz, George Verghese and Bernard Lesieutre, "Observer-Based Monitors for Electromechanical Dynamics in Power Networks," Proc. 15th Power System Computation Conference, August 2005, Liege, Belgium.
- [116] Teruo Ono, Game Theoretic Analysis and Agent-Based Simulation of Electricity Markets, Masters thesis, EECS Department, May 2005.
- [117] Teruo Ono and George Verghese, "Replicator Agents for Electricity Markets," 13th International Conf. on Intelligent Systems Application to Power Systems (ISAP 2005), Washington DC, November 2005.
- [118] M. Johansson, L. Xiao, and S. Boyd. Simultaneous routing and resource allocation in CDMA wireless data networks. Submitted to *The 2003 IEEE International Conference on Communications*, August 2002.
- [119] S. Lall and C. Beck. Error bounds for balanced model reduction of linear time-varying systems. *IEEE Transactions on Automatic Control*, 48(6):946–956, June 2003.
- [120] L. Xiao, M. Johansson, and S. Boyd. Simultaneous routing and resource allocation via dual decomposition. Submitted to *IEEE Transactions on Communications*, August 2002.
- [121] M. Rotkowitz and S. Lall. Decentralized control information structures preserved under feedback. In *Proceedings of the IEEE Conference on Decision and Control*, pages 569–575, December 2002.
- [122] R. Cogill and S. Lall. Topology independent controller design for networked systems. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 1788–1793, 2004.
- [123] M. Rotkowitz and S. Lall. On computation of optimal controllers subject to quadratically invariant sparsity constraints. In *Proceedings of the American Control Conference (ACC)*, pages 5659–5664, June 2004.
- [124] B.-D. Chen and S. Lall. Control of distributed discrete-time systems on graphs. In *Proceedings of the American Control Conference (ACC)*, pages 2251–2256, June 2004.

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

10 APR 2007

DTIC Data

Page 1 of 1

Purchase Request Number: FQ8671-0700076
BNP: F1ATA06202B0AC
Proposal Number: 05-NL-383
Research Title: A MATHEMATICAL MODEL OF THE HUMAN SLEEP-WAKE CYCLE
Type Submission: Annual Report #1
Inst. Control Number: FA9550-06-1-0033P00002
Institution: THE OHIO STATE UNIV RESEARCH FOUNDATION
Primary Investigator: Prof David Terman
Invention Ind: none
Project/Task: 2313E / X
Program Manager: Dr. Willard D. Larkin

Objective:

This research will develop a new, physiologically-informed mathematical model for the human sleep-wake cycle, aimed at a deeper understanding of human performance under adverse schedules of work and rest. The primary objective is to develop a valid scientific capacity to understand how sleep disruption and extended wakefulness interact with homeostatic and circadian mechanisms of sleep/wake control, and thus enable effective countermeasures for fatigue-induced deficits in human performance.

Approach:

Model development will proceed in stages. In the first stage, sleep-promoting cells within the hypothalamus, wake-promoting monoaminergic cell groups, and the circadian oscillator in the suprachiasmatic nucleus will be considered. Each cell will be modeled using the Hodgkin-Huxley formalism. In the second stage, other cell groups that have a role in ultradian rhythms and EEG activity will be added; transitions between sleep and wake and between REM and non-REM sleep will be investigated. In the third stage, focus will be on individual variability and on hypotheses for effective interventions for robust human performance. Data from research on human performance and on sleep neurobiology will be used to refine and validate the mathematics.

Progress:

Year: 2006 **Month:** 09

Not required at this time.

Year: 2007 **Month:** 04

ANNUAL REPORT #1 - FA9550-06-1-0033:

The primary aim of the research described in this proposal is to construct a detailed, biologically-based model for the human sleep-wake cycle. The goal is to design a model that more accurately predicts sleep-wake behavior and its relation to alertness, performance and fatigue. The biological basis of sleep and waking is extremely complex and many details of this system are not yet known. The model we are developing is being used to test hypotheses for mechanisms underlying the generation of this complex dynamics. The model has already led, for example, to a deeper understanding of how the circadian and homeostatic processes interact to generate the switch between sleep and wakefulness as well as between REM and non-REM sleep. We have already developed a model for the development of the circuitry controlling sleep. During the coming year, we will use these models to formulate and test hypotheses concerning the impact of sleep loss on human performance and protective interventions. We will also explore how individual differences should be incorporated into the model.